

WHAT IS CLAIMED IS:

- 1 1. A method comprising:
 - 2 for a first thread, entering a processing queue for obtaining permission to
 - 3 enter a critical section of code;
 - 4 determining if a second thread exists, the second thread executing the
 - 5 critical section of code concurrently with the first thread entering the
 - 6 processing queue; and
 - 7 if the second thread exists, then determining if the second thread is
 - 8 executing the critical section;
 - 9 if the second thread is executing the critical section, then testing for the
 - 10 second thread to complete until one of the following occurrences:
 - 11 the second thread completes; and
 - 12 the yielding count expires.
- 1 2. The method of claim 1, additionally comprising if the yielding count expires
- 2 before the second thread completes, then exiting the processing queue.
- 1 3. The method of claim 2 additionally comprising re-entering the processing
- 2 queue after a period of time.
- 1 4. The method of claim 3, wherein the period of time is determined by an
- 2 operating system scheduling algorithm.
- 1 5. The method of claim 1, additionally comprising if the second thread
- 2 completes before the yielding count expires, then executing the first critical
- 3 section of code.
- 1 6. The method of claim 1, additionally comprising if the second thread does
- 2 not exist, then executing the first critical section of code.
- 1 7. The method of claim 1, wherein the yielding count is based on the number

2 of threads contending to enter a corresponding critical section of code.

1 8. The method of claim 7, wherein the yielding count is based on twice the
2 number of threads contending for the lock.

1 9. The method of claim 1, wherein the yielding count is based on the number
2 of CPUs (central processing units).

1 10. The method of claim 1, wherein the critical section of code includes the
2 same code in both the first and the second thread.

1 11. A method comprising:

2 for a first thread, entering a processing queue for obtaining a lock on a
3 shared resource in a first critical section of code by checking the
4 status of shared variables existing in a memory, the shared
5 variables including a turn variable and a status flag;

6 determining if a second thread exists, the second thread executing a
7 second critical section of code concurrently with the first thread
8 entering the processing queue, the second critical section
9 corresponding to the second thread; and

10 if the second thread exists, then testing for the second thread to relinquish
11 the lock on the shared resource by testing the status flag, the
12 testing to be performed until one of the following occurrences:

13 the second thread relinquishes the lock when the flag has been
14 reset; and

15 the yielding count expires.

1 12. The method of claim 11, wherein if the second thread relinquishes the lock
2 before the yielding count expires, then obtaining the lock on the shared
3 resource.

- 1 13. The method of claim 11, wherein if the yielding count expires before the
2 second thread completes, then exiting the processing queue.
 - 1 14. The method of claim 13, additionally comprising re-entering the
2 processing queue after a determined amount of time.
 - 1 15. A method comprising:
 - 2 a. initializing shared variables, the shared variables including a
3 turn variable, a first status flag, and a second status flag;
 - 4 b. reading the shared variables into a memory;
 - 5 c. entering a processing queue;
 - 6 d. determining if a yield count has expired;
 - 7 e. if the yield count has expired, then exiting the processing queue;
 - 8 f. if the yield count has not expired, then for a contending process,
9 determining if a concurrent process exists;
 - 10 g. retrieving the second status flag and the turn variable from the
11 memory, reading the status flag into a first cache and reading
12 the turn variable into the second cache to determine if the
13 concurrent process is executing a critical section of code;
 - 14 h. if the concurrent process is not executing a critical section of
15 code, then entering the critical section of code, and upon
16 completing the critical section of code, resetting the first status
17 flag; and;
 - 18 i. if the concurrent process is executing the critical section of
19 code, then repeating d through i.

Page 10 of 10

- 1 16. The method of claim 15, wherein the second cache is larger than the first
2 cache.
- 1 17. The method of claim 15, wherein resetting the first status flag comprises
2 retrieving a reset value from a register.
- 1 18. A machine-readable medium having stored thereon data representing
2 sequences of instructions, the sequences of instructions which, when
3 executed by a processor, cause the processor to perform the following:
 - 1 for a first thread, enter a processing queue for obtaining permission to
2 enter a critical section of code;
 - 3 determine if a second thread exists, the second thread executing the
4 critical section of code concurrently with the first thread entering the
5 processing queue; and
 - 6 if the second thread exists, then determine if the second thread is
7 executing the critical section;
 - 8 if the second thread is executing the critical section, then test for the
9 second thread to complete until one of the following occurrences:
 - 10 the second thread completes; and
 - 11 the yielding count expires.
- 1 19. The method of claim 18, additionally comprising if the yielding count
2 expires before the second thread completes, then exiting the processing
3 queue.
- 1 20. The method of claim 18, additionally comprising if the second thread does
2 not exist, then executing the first critical section of code.
- 1 21. An apparatus comprising:
 - 2 at least one processor;

a machine-readable medium having instructions encoded thereon, which when executed by the processor, are capable of directing the processor to:

for a first thread, enter a processing queue for obtaining permission
to enter a critical section of code;

determine if a second thread exists, the second thread executing the critical section of code concurrently with the first thread entering the processing queue; and

if the second thread exists, then test for the second thread to complete until one of the following occurrences:

the second thread completes; and

the yielding count expires.

22. The method of claim 21, additionally comprising if the yielding count expires before the second thread completes, then exiting the processing queue.
 23. The method of claim 21, additionally comprising if the second thread does not exist, then executing the first critical section of code.
 24. An apparatus comprising:
 - means for a first thread to enter a processing queue for obtaining permission to enter a critical section of code;
 - means to determine if a second thread exists, the second thread executing the critical section of code concurrently with the first thread entering the processing queue; and
 - if the second thread exists, then means to determine if the second thread is executing the critical section;

9 if the second thread is executing the critical section, then means to test for
10 the second thread to complete until one of the following
11 occurrences:
12 the second thread completes; and
13 the yielding count expires.